



AFRL-RI-RS-TR-2013-041

CYBER ADVERSARY DYNAMICS

DARTMOUTH COLLEGE

FEBRUARY 2013

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2013-041 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ S /

RICHARD FEDORS
Work Unit Manager

/ S /

JULIE BRICHACEK
Chief, Information Systems Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) FEBRUARY 2013		2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) SEP 2011 – SEP 2012	
4. TITLE AND SUBTITLE CYBER ADVERSARY DYNAMICS				5a. CONTRACT NUMBER FA8750-11-1-0253	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER 61101E	
6. AUTHOR(S) GEORGE CYBENKO				5d. PROJECT NUMBER C2DY	
				5e. TASK NUMBER NM	
				5f. WORK UNIT NUMBER CS	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Dartmouth College Thayer School of Engineering 8000 Cummings Hall Hanover NH 03755				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RISC 525 Brooks Road Rome NY 13441-4505 DARPA/I20 675 N. Randolph St Arlington, VA 22203-2114				10. SPONSOR/MONITOR'S ACRONYM(S) N/A	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2013-041	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. PA# 88ABW-2013-0616 Date Cleared: 11 FEB 2013					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The goal of this project was to develop and demonstrate capabilities for modeling and exploiting the coevolution of offensive and defensive cyber behavior. We are calling such capabilities Cyber Adversary Dynamics. Using recent advances in behavioral game theory and a systematic treatment of open source data, this project created a scientific foundation for modeling cyber activity within adversarial situations. Initial research in this field has produced proof-of-concept approaches for and examples of Cyber Adversary Dynamics, identifying what is possible over different time scales and investment levels. This forms a basis for understanding and assessing adversarial cyber behavior, and provides a significant national opportunity for improving cyber resiliency.					
15. SUBJECT TERMS Adversary modeling; inverse game theory; behavioral economics; representing, anticipating, visualizing cyber behavior; cyber utilities, cyber operations, command and control.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON RICHARD FEDORS
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) NA

TABLE OF CONTENTS

Section.....	Page
List of Figures	ii
List of Tables.....	iii
1.0 SUMMARY	1
2.0 INTRODUCTION.....	2
3.0 METHODS, ASSUMPTIONS AND PROCEDURES	4
4.0 RESULTS AND DISCUSSION	8
4.1 Attacking and Defending Adversarial Behavioral Models	8
4.2 Attack Graph Games and their Asymptotic Equilibria	8
4.3 Botnet Design for Improved Stealth and Mission Effectiveness	8
4.4 An Analytic Approach to Cyber Adversarial Dynamics.....	9
4.5 Tools for Behavioral Anomaly Detection and Analysis	9
5.0 CONCLUSIONS AND RECOMMENDATIONS.....	11
6.0 REFERENCES	12
APPENDIX A – Publications and Presentations	13
APPENDIX B – Abstracts	14
LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS	18

List of Figures

Figure 1: Transportation fatalities rates between 1950 and 2003	4
Figure 2: Performance in real adversarial environments.....	6

List of Tables

Table 1: Scientific paradigms applicable to cyber security.....	2
Table 2: Variants of Game Theory and their applicability to Cyber C4I.....	2

1.0 SUMMARY

The goal of this project has been to develop and demonstrate capabilities for modeling and exploiting the co-evolution of offensive and defensive cyber technologies and tactics. We are calling such capabilities *Cyber Adversary Dynamics*. Dominance in cyber conflict will depend on our ability to recognize and eventually anticipate changes in adversaries' tactics, techniques and procedures (TTP's) in cyber operations.

To date, cyber security investment by both the government and commercial sectors have been largely driven by the "best response" to whatever threat environment was being perceived at the time. Adversaries see those investments and shift their attacks to less protected surfaces, software systems or parts of the supply chain while respecting their own resource constraints (skills, funding, time) and strategic needs. There has been appropriate hand wringing and complaining about this situation but this is the best anyone has to date.

This project will leverage recent advances in behavioral game theory, behavioral economics and adversary modeling by applying them to the more than twenty years of documented experience we now have in cyber security. We have much open source data about a.) cyber vulnerabilities; b.) commercial software releases, patches and updates; c.) exploits and; d.) actual attacks. Viewing this data in a principled, scientific manner will allow us to model the dynamics of cyber TTP's in a way that has not previously been done.

The long-range vision of this project is to ultimately achieve the kinds of capabilities enabled by DARPA programs such as Wide Area Network Detection and Insight, but for cyber operations. What makes them feasible today are the understanding, modeling and exploitation of adversaries' physical dynamics and organizational business processes. Such understanding, modeling and exploitation in the cyber domain do not presently exist.

This project strives to show how those foundations can be built. Success will improve all aspects of U.S. military and civilian cyber operations by providing tools for a strategic view of how threats and responses co-evolve, thereby better informing tactical operations. This is especially important for dealing with Advanced Persistent Threats because those campaigns are long term, strategic and costly to create and defend against.

This one-year project produced proof-of-concept approaches for and examples of Cyber Adversary Dynamics, identifying what is possible over different time scales and investment levels. We believe this small-scale project can identify a significant national opportunity for improving cyber resiliency.

A variety of scientific paradigms can be brought to bear on adversarial cyber security and behavioral science. These paradigms are enumerated in Table 1. This enumeration of paradigms, due to the Principal Investigator George Cybenko, has been disseminated in numerous presentations as well.

Table 1: Scientific paradigms applicable to adversarial cyber security.

Formalism	Primitives	Examples
Aristotelian	Objects, properties and relationships	Formal methods, expert systems, etc.
Newtonian	System state and dynamics	Control theory, Operations Research – deterministic and stochastic
Darwinian	Competition and constraints	Game theory, utility theory, pursuit and evasion, economics
Kahnemanian	Human decision making and behaviors	“Law of Small Numbers”, non-optimal behaviors, biases, priming, prospect theory

The present project has focused on the last two paradigms, specifically addressing human-cyber aspects of cyber operations. We have used a previously developed informal survey of adversarial modeling approaches in the context of different types of game theory resulting in the assessment summarized in Table 2.

Table 2. Variants of Game Theory and their applicability to Cyber Adversarial Dynamics

	Maturity	Applicability, Realism	Robustness	Preliminary Results
Games with Complete Information	high	low	low	low
Games with Incomplete or Imperfect State Information	high	high	medium	medium
Games with Incomplete or Imperfect Objective Information	high	medium	medium	low
Adaptation, Evolution, Learning in Games	medium	high	medium	medium
Hypergame Theory	medium	high	high	high
Behavioral Models	medium	high	medium	medium

Our conclusion based on these surveys was that building realistic models of cyber adversary behaviors and learning adversaries’ utilities from those behaviors in an actual operation (or game play) was an important area for immediate investigation.

That work has been documented in detail in several publications resulting from this research, including:

Sweeney, P.T., "Botnet Design for Improved Stealth and Mission Effectiveness," *Ph.D. thesis proposal, Dartmouth College*, 2012.

Sweeney, P., & Cybenko, G. "An analytic approach to cyber adversarial dynamics." In *SPIE Defense, Security, and Sensing* (pp. 835906-835906). International Society for Optics and Photonics. (2012, May).

Cybenko, G., and Landwehr, C. E. "Security Analytics and Measurements." *IEEE Security & Privacy*, 10(3), 5-8. (2012).

Stocco, G. and Cybenko, G., "Exploiting Adversary's Risk Profiles in Imperfect Information Security Games," *Proceedings of Conference on Decision and Game Theory for Security (GAMESEC 2011)*, November 14-15, 2011, College Park, Maryland, USA.

Bilar, D., Cybenko, G. and Murphy, J. "Adversarial Dynamics: The Conficker Case Study." *Moving Target Defense II* (2012): 41-71.

Stocco, G.F., "Toward a precise simulation model for multi-agent computer security scenarios", *Ph.D. thesis proposal, Dartmouth College*, 2012.

Cybenko, G. and Stocco, G.F., "Attack Graph Games and their Asymptotic Equilibria," in preparation.

Berk, V. H., Cybenko, G., Souza, I. G. D., & Murphy, J. P. (2012, January). "Managing Malicious Insider Risk through BANDIT." In *System Science (HICSS), 2012 45th Hawaii International Conference on* (pp. 2422-2430). IEEE.

Research documented in these publications is summarized in this report.

3.0 METHODS, ASSUMPTIONS AND PROCEDURES

Background

Progress in operational cyber security remains difficult to demonstrate. In spite of the research and development investments made over more than 30 years, many government, commercial and consumer information systems continue to be successfully attacked and exploited on a routine basis. By contrast, research and development investments in, for example, automobile and aviation safety over the same time periods have led to significant, demonstrable improvements in the corresponding domains. Figure 1 illustrates how transportation fatality rates have dropped between 1950 and 2003 for passenger vehicles, commercial air travel, and passenger rail service. It is clear that investments in transportation safety have resulted in quantifiable progress using these metrics. The same cannot be said, yet, of investments in cyber security.




Figure 1. Transportation fatality rates between 1950 and 2003 for passenger vehicles (left), commercial air travel (middle), and passenger rail (right).

[Sources: National Highway Traffic Safety Administration; Air Transport Association; Federal Railroad Administration, Bureau of Transportation Statistics]

A major difference between automobile safety and information security is that in the former the “adversaries” are primarily natural laws that don’t change while in the latter the adversaries are rational humans who adapt quickly and creatively. Consequently, we cannot understand or model the cyber security landscape in terms of steadily making progress towards “a solution.”

We must view cyber security as an ongoing sequence of moves, countermoves, deceptions and strategic adaptations by the various actors involved – attackers, defenders, vendors and decision/policy makers. Accordingly, we believe that the appropriate science for understanding the evolving landscape of cyber security cannot be found in the logic of formal systems or new software engineering techniques; it is an emerging subarea of game theory that investigates *dynamics* in adversarial situations and the biases of competing human agents that drive those dynamics.

Adversarial Dynamics vs Classical Game Theory

The original goals of Game Theory were to model adversarial environments and to optimize strategies for operating in those environments. This would seem ideal for modeling cyber operations as well as other national security situations - indeed there is a community of researchers currently investigating the application of classical Game Theory to information assurance and cyber operations.

However, the overwhelming focus of Game Theory research over the past 60 years has been on the problem of “solving” games that are defined a priori. That is, most Game Theory research to date begins by assuming a game is already defined (namely, the players, their possible moves and payoffs) and then explores properties of optimal strategies and how to compute them. Optimality is with respect to a solution criterion such as Nash Equilibrium or Pareto Optimality.

An obvious criticism of this approach is that in most real world adversarial situations we do not know who the players are, what their possible moves might be and, perhaps most importantly, what their preferred outcomes or objectives are. Put another way, none of the players actually know the complete details of the game that they are playing! A further complication is that few people outside of the Game Theory cognoscenti know what a Nash Equilibrium is, let alone how to compute one (which turns out to be quite difficult for real, complex multi-player situations) and how to implement the strategy in actual play.

As a result, while Game Theory can inform us about how to play chess, checkers, poker and the simple examples found in most Game Theory texts, it has not been as useful in the majority of real-world adversarial situations. New directions and ideas are needed, especially in the area of cyber security.

Adversarial Dynamics

Adversarial behavior dynamics is the empirical study of players’ actions in adversarial situations. The “game” in these adversarial situations is implicit and can only be understood in terms of the moves players make and how they evolve their play in response to each other’s moves. Figure 2 illustrates such evolving behavior as oscillatory “performance,” which characterizes activity in real adversarial environments. The top left graph shows the number of vulnerabilities and exploits reported in OSVDB over time. On the top right those exploits have been normalized per billion dollars of e-commerce. The bottom left graph shows the number of apprehensions by the United States Border Patrol along the southwest U.S. border. The bottom right graph shows the first-offense types for juveniles in California over time. Our assertion is that oscillations such as these are not only typical and supported by data but are intrinsic to adversarial dynamics in many domains, including cyber operations. The challenge for us is to model, learn, and ultimately exploit such oscillations, especially in cyber operations.

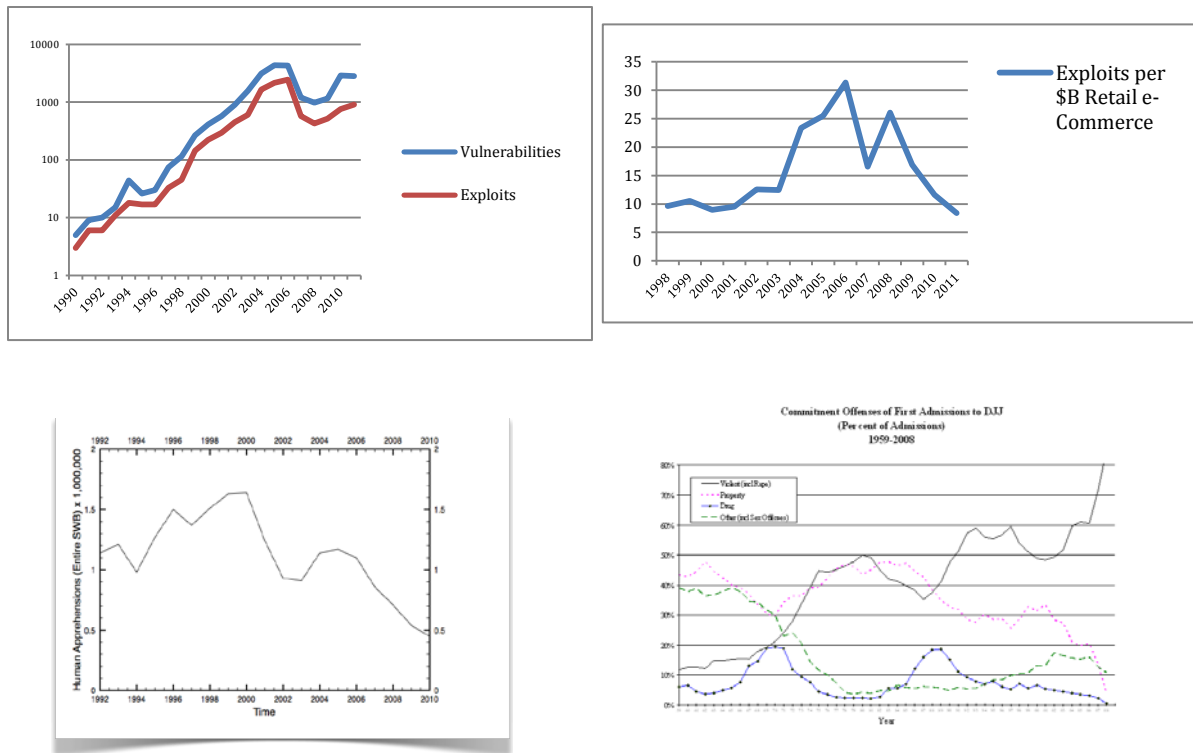


Figure 2. Oscillatory performance in real adversarial environments.
 (Top left: number of vulnerabilities and exploits reported in OSVDB over time. Top right: exploits normalized per billion dollars of e-commerce. Bottom left: number of apprehensions by the United States Border Patrol along the U.S. southwest border. Bottom right: first-offense types for juveniles in California over time in years.)

Our research group has studied historical data from a variety of cyber and national security domains such as computer vulnerability databases, offensive and defense coevolution of wormbots such as Conficker, and US border security. The data show that the “success rates” in these different domains oscillate over time – they are not converging to any asymptote. Such oscillations are indicators of and intrinsic to adversarial dynamics in complex, competitive environments. In particular, each player is adapting incrementally to the observed play of their opponents. This can be modeled by systems of differential equations known as “replicator equations.”

The replicator equations are typically nonlinear (third order nonlinear, in fact) so that the resulting dynamics are difficult to predict analytically. However, the inverse problem of observing behaviors and estimating parameters of the replicator equations that result in those behaviors are tractable computational problems. In particular, it is possible to observe game play and strategy evolution and then make inferences about the players’ motives, costs and move options.

This kind of modeling approach can explain the non-convergent dynamics we are seeing in cyber security. Moreover, it seems to be a necessary ingredient to analyzing “big data” cyber security problems. Recognizing such adversarial dynamics over the full spectrum of information operations can give us a handle on what is really going on and can help us forecast the various players’ future moves and strategies. Recognizing and

harnessing the realities of such dynamic coevolution will be a key ingredient to dominating cyber operations.

Implications for Military Cyber Operations

Technologies and tactics in new conflicts are often surprising, despite the documented history of three thousand years of kinetic warfare available to us. While humans adapt by quickly learning from previous experience, especially so in life threatening situations, we are still not very good at anticipating how others will react in novel open-ended, competitive environments.

The situation with respect to cyber conflict is even worse than kinetic warfare. We have very limited experience in real cyber conflict, let alone “total cyber war” with peer or near-peer adversaries who are capable of advanced, combined kinetic and cyber operations in which civilian and industrial infrastructures are targeted. Compounding matters is the fact that the underlying technologies (software, hardware and communications systems) are constantly changing and poorly understood, by their developers and users alike.

A true cyber conflict will be more novel and surprising than any previous conflict unless we deliberately and creatively prepare better. A key part of such preparation is an understanding of the different agents’ utilities, capabilities, tactics and strategies in a co-evolving context.

The different agents involved in cyber operations include but are limited to: defenders, attackers (a variety of state sponsored, organized non-state actors, ideological groups and cyber mercenaries for example), commercial vendors, system administrators, policy makers and end users. All such groups contribute to the current operational cyber environment in different ways.

- Each actor typically has different: *Utilities* – preferences between different states of the world that they can bring about;
- *Capabilities* – the sets of actions available to an actor as determined by technical prowess, legal constraints, finances, human resources for example;
- *Tactics* – Short term goals and techniques to achieve those goals;
- *Strategy* – Long term outcomes and techniques to achieve those outcomes.

This project surveyed the use of historical time-stamped data about vulnerabilities, exploits, defensive technologies and actual incidents. A large corpus of open source data on these ingredients already exists. We are aware that proprietary repositories and histories of malware also exist within the government and commercial sectors but we will demonstrate the key concepts on open source information solely. Techniques developed will be applicable to other sources as well.

4.0 RESULTS AND DISCUSSION

4.1 Attacking and Defending Adversarial Behavioral Models

In this research, we developed techniques for attacking and defending behavioral anomaly detection methods commonly used in network traffic analysis and covert channels. The main new result is a demonstration of how to use an adversary's dynamical behavior's k -order statistics to build a stochastic process that has the same k -order stationary statistics but possesses different, deliberately designed, $(k+1)$ -order statistics if desired. Such a model realizes a "complexification" of the process or behavior that a defender can use to monitor whether an attacking adversary is shaping the behavior. We also describe a source coding technique that respects the k -order statistics, including entropy, which is a first order statistic for example, of a process while encoding information covertly. Although the main results and examples are stated in terms of behavioral anomaly detection for covert channels, the techniques are more generally applicable to behavioral anomaly analysis. One fundamental consequence of these results is that certain types of adversarial behavioral anomaly detection techniques come down to an arms race in the sense that the advantage goes to the party that has more computing resources applied to the problem.

4.2 Attack Graph Games and their Asymptotic Equilibria

This work proposes a quantitative formulation of attack graph optimization suitable for modeling certain adversarial cyber attack/defend scenarios. The formulation is based on representing an attack as a finite graph in which directed edges represent the steps in an attack and edge weights representing the estimated costs to an attacker for traversing the edge. An attacker strives to traverse the graph from a specified start node to a specified end node using the shortest directed path between those nodes. On the other hand, the defender seeks to allocate defensive measures in such a way as to maximize the attacker's minimal cost attack path. We study the role that minimal cut sets play in hardening the attack graph and prove that minimal cut sets are optimal defensive investments in the limit although they may not play a role initially.

4.3 Botnet Design for Improved Stealth and Mission Effectiveness

Botnets (internet connected computers controlled by a single botmaster) present a great threat to Internet security. The moniker is almost universally associated with malicious and criminal activities, but that ignores the possibility to use botnets for legitimate operations in cyberspace as well. Regardless of a botnet's intended use, a review of literature on botnets turns up many references to *stealth*. Botnet *effectiveness* is discussed explicitly to a lesser extent, but of course a botnet must provide satisfactory mission effectiveness or the botmaster will change tactics. Both stealth and effectiveness are critically important to a botmaster, yet they're notoriously difficult to quantify.

This work has proposed a set of mission effectiveness and stealth measures that are derived from knowledge of the target network topology. Armed with such information, the botmaster can design a botnet for interesting and subtle missions that go beyond what we've come to expect (e.g. distributed denial of service attacks that rely primarily on a botnet's

massive scale). By understanding and defining specific effectiveness and stealth objectives for a known network topology, the botmaster can quantify how effective and stealthy the botnet will be during its mission. Advancing these measures one step further, our goal is to provide a botnet design engine that will, given a set of effectiveness and stealth objectives, enhance the capability to identify the *cyber high ground*—the set of systems that, if controlled by the botmaster, yield the highest probability of mission success. Initial results show that mission effectiveness and stealth can be defined and measured through a set of basic objectives, and that capability can enable design of botnets that are best suited to the environments and mission for which they will be employed.

4.4 An Analytic Approach to Cyber Adversarial Dynamics

To date, cyber security investment by both the government and commercial sectors has been largely driven by the myopic best response of players to the actions of their adversaries and their perception of the adversarial environment. However, current work in applying traditional game theory to cyber operations typically assumes that games exist with prescribed moves, strategies, and payoffs. This report presents an analytic approach to characterizing the more realistic cyber adversarial meta-game that we believe is being played. Examples show that understanding the dynamic meta-game provides opportunities to exploit an adversary's anticipated attack strategy. A dynamic version of a graph-based attack-defend game is introduced, and a simulation shows how an optimal strategy can be selected for success in the dynamic environment.

4.5 Tools for Behavioral Anomaly Detection and Analysis

As our communication infrastructure has become more ubiquitous, we have shifted the implementation of our business processes increasingly to this information infrastructure. Things that in the past required a phone call are now available over email, chat, or solid media. Documents have moved from filing cabinets to file shares, customer rolodexes to databases, and accounting books to spreadsheets. A filing cabinet, however, is often much less accessible than a networked server.

The network security industry has therefore been profiting handsomely from the sale of access control and security solutions, which work very well as long as an insider is trustworthy. Document-level access control solutions, such as the various Digital Rights Management (DRM) schemes do not mitigate the threat posed by trusted insiders. Although fine-grained access control to information helps to some degree, any person with qualified access and malicious intentions still has the power to steal or destroy the information.

True protection from the threat of a malicious, trusted insider is impossible to achieve, but traditional ideas implemented with a modern twist give us some insight into the problem, and help us minimize the risk. Using modern behavioral analysis techniques, combined with measured behavioral baselines, we score each user's actions for their means, their perceived motive, and their measured opportunity for committing malicious insider actions (the MMO-score). Users' behavior is compared both to their own baselines, as well as the behaviors of members in their group, for all three axes of the MMO-score. This creates a comprehensive 3-dimensional view of 'abnormality' for each user.

The system is designed to focus an investigative analyst on those users that show the greatest deviation and risk on all three axes of the MMO-score, without giving a conclusive yes/no decision on malicious intent automatically. By organizing an often-bewildering array of approaches to the insider threat problem into this conceptual framework, the problem of evaluating multiple streams of data becomes more tractable and intuitive. This effectively forms a powerful insider threat decision support system.

Moreover, we have developed a general quantitative analysis technique for assessing false positive and false negative rates in such detection systems.

5.0 CONCLUSIONS AND RECOMMENDATIONS

This research has investigated a variety of game theoretic and adversarial reasoning approaches to modeling cyber adversarial behaviors. Among our key findings are that many cyber security techniques do not “look ahead” at their consequences with respect to how and when an adversary will adapt as a result of the technology.

Approaches to anticipating adversarial covert channel manipulations we have developed include a new and powerful algorithmic technique based on principled stochastic modeling theory for both the defenders and attackers of covert channels. This effectively reduces the operational use of covert channels to an arms race of computing and memory resources.

We have identified a variety of approaches to defining the “high ground” in a computer network. By analogy with classical, physical space military doctrine and practice, certain parts of a network are advantageous for a combatant to “own.” We have developed a framework in which these ideas can be related to stealth and mission specifics. Moreover, these quantifications are computable and operationally meaningful to botnet operators in particular.

These lines of research appear to be very promising for future investigation.

Recommended directions include but are not limited to:

- Design and implement a continuously updated database of government and military network security incidents which can be correlated to vulnerability and malware databases so that adversarial dynamics trends can be estimated dynamically and in real time;
- Build a similar database for detected information exfiltration and covert channels;
- Design and develop interactive tools for exploring the above databases so that U.S. government and military cyber defenders can quickly identify and explore adversarial dynamics from observed incidents.

REFERENCES

- Bilar, D., Cybenko, G. and Murphy, J. "Adversarial Dynamics: The Conficker Case Study." *Moving Target Defense II* (2012): 41-71.
- Gupta, A.; Kuppili, P.; Akella, A.; Barford, P. "An empirical study of malware evolution," *Communication Systems and Networks and Workshops, 2009. COMSNETS*. pp.1-10, Jan. 2009. <http://pages.cs.wisc.edu/~archit/projects/malware/>
- National Vulnerability Database Version 2.2, <http://nvd.nist.gov/>
- The Open Source Vulnerability Database, <http://osvdb.org/>
- Comparison of Microsoft Windows versions, <http://osvdb.org/>
- Timeline of viruses and worms,
http://en.wikipedia.org/wiki/Timeline_of_computer_viruses_and_worms
- Malware History,
http://download.bitdefender.com/resources/files/Main/file/Malware_History.pdf
- <http://www.levenez.com/windows/>
- House, J.T., **Game-Theoretic Approaches for Adversarial Multi-Armed Bandit Scenarios**. *Ph.D. Thesis, Dartmouth*, 2012.
- House, J.T. and Cybenko, G. "Hypergame theory applied to cyber attack and defense." *SPIE Defense, Security, and Sensing. International Society for Optics and Photonics*, 2010.
- House, J.T. and Cybenko, G. "Exploiting exploration strategies in repeated normal form security games." *SPIE Defense, Security, and Sensing. International Society for Optics and Photonics*, 2012. Kocsis, L. and Szepesvari C., "Bandit based Monte-Carlo planning," *Machine Learning: ECML 2006* **4212**, pp. 282–293, 2006.
- Lai, T. and Robbins, H., "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics* **6**, pp. 4–22, 1985.
- Landwehr, C., "Cybersecurity: From engineering to science," *National Security Agency The Next Wave*, **Vol. 19**, No. 2, 2012.
- Stocco, G.F., **Toward a precise simulation model for multi-agent computer security scenarios**, *Ph.D. thesis proposal, Dartmouth College*, 2012.
- Stocco, G. and Cybenko, G., "Exploiting Adversary's Risk Profiles in Imperfect Information Security Games," *Proceedings of Conference on Decision and Game Theory for Security (GAMESEC 2011)*, November 14-15, 2011, College Park, MD.

APPENDIX A: Publications and Presentations

Sweeney, P.T., "Botnet Design for Improved Stealth and Mission Effectiveness," *Ph.D. thesis proposal, Dartmouth College*, 2012.

Sweeney, P., & Cybenko, G. "An analytic approach to cyber adversarial dynamics." In *SPIE Defense, Security, and Sensing* (pp. 835906-835906). International Society for Optics and Photonics. (2012, May).

Cybenko, G., and Landwehr, C. E. "Security Analytics and Measurements." *IEEE Security & Privacy*, 10(3), 5-8. (2012).

Stocco, G. and Cybenko, G., "Exploiting Adversary's Risk Profiles in Imperfect Information Security Games," *Proceedings of Conference on Decision and Game Theory for Security (GAMESEC 2011)*, November 14-15, 2011, College Park, Maryland, USA.

Bilar, D., Cybenko, G. and Murphy, J. "Adversarial Dynamics: The Conficker Case Study." *Moving Target Defense II* (2012): 41-71.

Stocco, G.F., "Toward a precise simulation model for multi-agent computer security scenarios", *Ph.D. thesis proposal, Dartmouth College*, 2012.

Cybenko, G. and Stocco, G.F., "Attack Graph Games and their Asymptotic Equilibria," in preparation.

Berk, V. H., Cybenko, G., Souza, I. G. D., & Murphy, J. P. (2012, January). "Managing Malicious Insider Risk through BANDIT." In *System Science (HICSS), 2012 45th Hawaii International Conference on* (pp. 2422-2430). IEEE.

APPENDIX B: Abstracts

Bilar, D., Cybenko, G. and Murphy, J. "Adversarial Dynamics: The Conficker Case Study." *Moving Target Defense II* (2012): 41-71.

Abstract: It is well known that computer and network security is an adversarial challenge. Attackers develop exploits and defenders respond to them through updates, service packs or other defensive measures. In non-adversarial situations, such as automobile safety, advances on one side are not countered by the other side and so progress can be demonstrated over time. In adversarial situations, advances by one side are countered by the other and so oscillatory performance typically emerges. This paper contains a detailed study of the coevolution of the Conficker Worm and associated defenses against it. It demonstrates, in concrete terms, that attackers and defenders each present moving targets to the other. After detailing specific adaptations of attackers and defenders in the context of Conficker and its variants, we briefly develop a quantitative model for explaining the coevolution based on what we call Quantitative Attack Graphs (QAG) which involve attackers selecting shortest paths through an attack graph and defenders investing in hardening the shortest path edges appropriately.

Stocco, G.F., Toward a precise simulation model for multi-agent computer security scenarios, *Ph.D. thesis proposal, Dartmouth College*, 2012.

Abstract: Adversarial computer attacks are estimated to cost up to \$1 trillion per year. Aside from the raw fiscal costs of such compromises, there is a significant lack of understanding of how vulnerable systems really are, and further how much a compromise of such systems would really cost. Models exist for the modeling of these adversarial scenarios, but require a drastic simplification of the system at study, or a priori values not based on empirical data or real world observations. These solutions do not provide meaningful information to decision makers who are faced with the highly costly dangerous adversarial environment as it exists. This thesis proposes to create a novel adaptation of signaling games, a variant of Bayesian games as a model for real world computer security scenarios. This modified signaling game allows for more realistic modeling by using empirical data of real world scenarios and a more complex extensible game framework than previous, more restrictive models. In these scenarios the Defender is the operator of a system and the Adversary may either be an external or internal threat. While most of the discussion is framed in the context of a networked server the analysis holds as well for the operation of a single system where the Adversaries control processes while the Defender controls the operating system. Our goal is to create a continuous time multi-agent game theoretic model which will verifiably simulate the above outlined real world scenario, and will allow a Defender to evaluate the effectiveness of their defense strategy based on their goals and utility function.

Stocco, G. and Cybenko, G., "Exploiting Adversary's Risk Profiles in Imperfect Information Security Games," *Proceedings of Conference on Decision and Game Theory for Security (GAMESEC 2011)*, November 14-15, 2011, College Park, Maryland, USA.

Abstract. At present much of the research that proposes to provide solutions to Imperfect Information Non Cooperative games provides superficial analysis which then requires a priori knowledge of the game to be played. We propose that High Card, a simple Multiplayer Imperfect Information Adversarial game, provides a more robust model for such games, and further, that these games may model situations of real world security and international interest. We have formulated two such real world models, and have created a modeling bot, which when facing adversaries with equal or better performing risk profiles, achieves a 7-fold increase in win performance.

Sweeney, P.T., "Botnet Design for Improved Stealth and Mission Effectiveness," *Ph.D. thesis proposal*, Dartmouth College, 2012.

Abstract: Botnets (internet connected computers controlled by a single botmaster) present a great threat to Internet security. The moniker is almost universally associated with malicious and criminal activities, but that ignores the possibility to use botnets for legitimate operations in cyberspace as well. Regardless of a botnet's intended use, a review of literature on botnets turns up many references to stealth. Botnet effectiveness is discussed explicitly to a lesser extent, but of course a botnet must provide satisfactory mission effectiveness or the botmaster will change tactics. Both stealth and effectiveness are critically important to a botmaster, yet they're notoriously difficult to quantify.

This work has proposed a set of mission effectiveness and stealth measures that are derived from knowledge of the target network topology. Armed with such information, the botmaster can design a botnet for interesting and subtle missions that go beyond what we've come to expect (e.g. distributed denial of service attacks that rely primarily on a botnet's massive scale). By understanding and defining specific effectiveness and stealth objectives for a known network topology, the botmaster can quantify how effective and stealthy the botnet will be during its mission. Advancing these measures one step further, our goal is to provide a botnet design engine that will, given a set of effectiveness and stealth objectives, enhance the capability to identify the cyber high ground—the set of systems that, if controlled by the botmaster, yield the highest probability of mission success. Initial results show that mission effectiveness and stealth can be defined and measured through a set of basic objectives, and that capability can enable design of botnets that are best suited to the environments and mission for which they will be employed.

Sweeney, P., & Cybenko, G. "An analytic approach to cyber adversarial dynamics." *In SPIE Defense, Security, and Sensing* (pp. 835906-835906). *International Society for Optics and Photonics*. (2012, May).

Abstract: To date, cyber security investment by both the government and commercial sectors has been largely driven by the myopic best response of players to the actions of their adversaries and their perception of the adversarial environment. However, current work in applying traditional game theory to cyber operations typically assumes that games exist with prescribed moves, strategies, and payoffs. This paper presents an analytic approach to characterizing the more realistic cyber adversarial metagame that we believe is being played. Examples show that understanding the dynamic metagame provides opportunities to exploit an adversary's anticipated attack strategy. A dynamic version of a graph-based attack-defend game is introduced, and a simulation shows how an optimal strategy can be selected for success in the dynamic environment.

Cybenko, G., and Landwehr, C. E. "Security Analytics and Measurements." *IEEE Security & Privacy*, 10(3), 5-8. (2012).

Abstract: We must view cyber security as an ongoing sequence of moves, countermoves, deceptions and strategic adaptations by the various actors involved – attackers, defenders, vendors and decision/policy makers. Accordingly, we believe that the appropriate science for understanding the evolving landscape of cyber security cannot be found in the logic of formal systems or new software engineering techniques; it is an emerging subarea of game theory that investigates dynamics in adversarial situations and the biases of competing human agents that drive those dynamics.

Cybenko, G. and Stocco, G.F., "Attack Graph Games and their Asymptotic Equilibria," in preparation.

Abstract: This work proposes a quantitative formulation of attack graph optimization suitable for modeling certain adversarial cyber attack/defend scenarios. The formulation is based on representing an attack as a finite graph in which directed edges represent the steps in an attack and edge weights representing the estimated costs to an attacker for traversing the edge. An attacker strives to traverse the graph from a specified start node to a specified end node using the shortest directed path between those nodes. On the other hand, the defender seeks to allocate defensive measures in such a way as to maximize the attacker's minimal cost attack path. We study the role that minimal cut sets play in hardening the attack graph and prove that minimal cut sets are optimal defensive investments in the limit although they may not play a role initially.

Berk, V. H., Cybenko, G., Souza, I. G. D., & Murphy, J. P. (2012, January). "Managing Malicious Insider Risk through BANDIT." *In System Science (HICSS), 2012 45th Hawaii International Conference on* (pp. 2422-2430). IEEE.

Abstract: The transition from system-to information-based security has continued steadily over the last 30 years. Correspondingly, it is increasingly not the computer that is at risk, but the information in it. The human operator is ultimately the

cornerstone of information security, an integral part of the information infrastructure. We are therefore forced to use techniques and methods that help us understand the role of human actors in the information infrastructure, so that we may make meaningful progress in mitigating insider threat. Malicious versus benign human behavior cannot easily be categorized based on a signature such as conventional virus and intrusion detection approaches. Because the cost of a false positive is high, we must be careful in our classification and subsequent actions. This article outlines our BANDIT (Behavioral Anomaly Detection for Insider Threat) system, using the traditional notion of Motive, Means, and Opportunity, combined with comprehensive behavioral analysis techniques to place each individual on a sliding scale of 'insider risk'. Finally, an insider threat detection cost-benefit analysis, based on classical risk assessment techniques, is presented to quantify how effective the technology has to be for beneficial deployment in a given enterprise.

LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

Quantitative Attack Graphs	QAG
Tactics, Techniques and Procedures	TTP
Defense Advanced Research Projects Agency	DARPA
Behavioral Anomaly Detection for Insider Threat	BANDIT
Means, Motive and Opportunity	MMO
Digital Rights Management	DRM